

Properties of spatial coupling in compressed sensing

Francesco Caltagirone

Institut de Physique Théorique
CEA Saclay and URA 2306, CNRS
91191 Gif-sur-Yvette, France.

Lenka Zdeborová

Institut de Physique Théorique
CEA Saclay and URA 2306, CNRS
91191 Gif-sur-Yvette, France.

Abstract—In this paper we address a series of open questions about the construction of spatially coupled measurement matrices in compressed sensing. For hardware implementations one is forced to depart from the limiting regime of parameters in which the proofs of the so-called threshold saturation work. We investigate quantitatively the behavior under finite coupling range, the dependence on the shape of the coupling interaction, and optimization of the so-called seed to minimize distance from optimality. Our analysis explains some of the properties observed empirically in previous works and provides new insight on spatially coupled compressed sensing.

I. INTRODUCTION

Spatial coupling is a methodological concept that permits to enhance considerably limits of algorithmic tractability in a wide range of inference problem where the underlying graphical model can be designed. Spatial coupling was originally introduced in low density parity check (LDPC) codes by [1], followed by a line of theoretical work on (terminated) convolutional LDPC codes, notably [2], [3]. The term itself, as well as a large part of the theoretical proofs and understanding comes from the work of [4]. The fact that with spatial coupling design of the system of interest one obtains tractable algorithms working down to the corresponding information theoretical limits (a phenomenon termed “threshold saturation”) has been successfully applied in a number of problems in error correcting codes, signal processing, communication systems and computer science problems, for an incomplete list of examples see references in [5].

One particularly important problem where spatial coupling was successfully applied is compressed sensing. It is well known that most signals of interest are compressible and this fact is widely used to save storage place. But usually compression is made only once the full signal is measured. In many applications (e.g. medical imaging using MRI or tomography) it is desirable to reduce the measurement time as much as possible (to reduce costs, radiation exposition dose etc.). Hence it is desirable to measure the signal directly in the compressed format. Compressed sensing is a concept implementing exactly this idea by designing new measurement protocols and sparse signal reconstruction algorithms [6], [7].

The first investigation of spatial coupling in compressed sensing is due to [8] where some limited improvement of performance was observed. A combination of Bayesian probabilistic approach, the use of approximate message passing algorithm and a special design of the (spatially coupled and seeded) measurement matrix led to an empirical evidence

and a conjecture that with spatial coupling the information theoretic limits can be indeed achieved in compressed sensing [9], [10]. This conjecture was proven rigorously by [11]. Recent generic and simple proof of the so-called threshold saturation also applies to the case of compressed sensing [12].

The view of possible experimental implementations of spatially coupled measurement matrices in compressed sensing devices (e.g. optical imagers using highly flexible arrays of micro-mirrors) motivates the need of more detailed study of this subject. The main open question that needs to be resolved concerns the optimization of the spatially coupled measurement matrix for systems of finite size. Such a process includes optimization of the number of blocks, range and shape of interactions. It also includes minimization of the “termination cost”, i.e. the determination of the size and strength of the seed that minimizes the final undersampling rate while assuring a robust propagation of the “successful reconstruction” wave. In this paper we address some of these questions.

II. SPATIAL COUPLING FOR COMPRESSED SENSING

In compressed sensing we consider a N -dimensional signal (vector) \mathbf{x} , this vector is K -sparse, i.e. only $K < N$ of its elements are nonzero. The measurements device records $M < N$ linear projections of this signal into a M -dimensional vector \mathbf{y} . These projections can be represented by a $M \times N$ measurement matrix F , such that $\mathbf{y} = F\mathbf{x} + \xi$ where ξ is (element-wise) an additive white Gaussian noise of variance Δ . The goal is to reconstruct tractably the sparse vector \mathbf{x} from the knowledge of \mathbf{y} and F for a given sparsity $\rho = K/N$ and the lowest possible undersampling ratio $\alpha = M/N$. Here we consider a model case of compressed sensing where the non-zero elements of the signal \mathbf{x} are taken as independent normally distributed variables of zero mean and unit variance, and we consider the asymptotic limit $N \rightarrow \infty$. We aim to minimize the mean squared error between the reconstructed and true signal $E = \sum_{i=1}^N (x_i^{\text{inferred}} - x_i^{\text{true}})^2 / N$.

To reconstruct the signal we work in the Bayesian setting where the goal is to sample the posterior probability distribution $P(\mathbf{x}|F, \mathbf{y}) = P(\mathbf{x})P(\mathbf{y}|F, \mathbf{x})/Z$. Such sampling is in general computationally difficult ($\#P$ -complete). An influential line of work [13], [14] introduced the approximate message passing (AMP), an iterative algorithm closely related to belief propagation (BP), as an algorithm that performs this sampling efficiently as long as the elements of the matrix F

are random independent variables and the system size N is large. Similarly to BP decoding in LDPC codes the AMP only works for some region of parameters ρ , α , Δ . In the regime of low measurement noise Δ there is an algorithmic threshold $\alpha_{\text{BP}}(\Delta, \rho)$ below which the algorithm fails, yet a good reconstruction of the signal is in principle possible down to an undersampling rate $\alpha_c(\Delta, \rho) < \alpha_{\text{BP}}(\Delta, \rho)$. The values of α_{BP} were obtained rigorously via state evolution asymptotic analysis of the AMP algorithm [15], the values α_c are known non-rigorously thanks to the replica calculation [9], [10].

The remarkable result of [9], [10], [11] is that when the concept of spatial coupling is used to design the measurement matrix F then the AMP algorithm reconstructs successfully down to measurement ratios α_c (this is analogous to the “threshold saturation” in LDPC codes).

The seeded spatially coupled matrix F used in this work is defined as follows, see Fig. 1. We split the N components of the signal \mathbf{x} into L equally sized blocks. We then split the components of the measurements \mathbf{y} into L blocks, each of the first w_s of them will have size $\alpha_s N/L$ (we call these first w_s blocks the “seed”), each of the remaining blocks will have size $\alpha_b N/L$ (we call these remaining blocks the “bulk”). We choose the elements of each of the block to be normally distributed with zero mean and variance J_{qr}/N , where $q, r = 1, \dots, L$. The coupling matrix J_{qr} is taken as

$$J_{qr} = \frac{J}{w} g\left(\frac{r-q}{w}\right), \quad \text{where} \quad \frac{1}{w} \sum_{z=-w}^w g\left(\frac{z}{w}\right) = 1, \\ g\left(\frac{z}{w}\right) = 0, \quad \text{for} \quad z \in (-\infty, -w) \cup (w, \infty). \quad (1)$$

This definition is very similar though not identical to the definitions in the previous literature on spatial coupling in compressed sensing. Note that [9] treated only the case of $w_s = w = 1$, in [10] generic w was considered and other designs of the matrix suggested, but without a systematic study of their advantages. The theoretical proofs of threshold saturation [11], [12] deal with the limit $1 \ll w \ll L \ll N$ and the seed is either not considered at all or has far too big cost for a practical implementation. In this limit the reconstructed part of the signal propagates as a kind of a travelling wave into the whole system [16]. Given that the system sizes of practical interest range from about $N = 10^3$ to say $N = 10^7$ (depending on the application) this limit cannot be implemented very closely. The main goal of this work is to still consider infinite block size, and to study the dependence on the values of parameters w , L , w_s , α_s and the function $g(\cdot)$ in order to achieve reconstruction of the signal under smallest possible average undersampling ratio

$$\alpha_{\text{eff}} = \frac{\alpha_b(L - w_s) + \alpha_s w_s}{L}. \quad (2)$$

When the size of every block is infinite the state evolution equations were derived in [9], [11] to describe the asymptotic




Fig. 1. A schematic representation of the seeded spatially coupled measurement matrix F . In the example above the number of blocks is $L = 10$, the interaction range is $w = 2$, the seed size is $w_s = 3$, the bulk and seed undersampling ratios are, respectively, $\alpha_b = 0.4$ and $\alpha_s = 0.8$. In light blue are the bulk blocks, in red the seeding blocks, in blue the interaction blocks, and in white the null elements.

evolution of the mean-squared error E_p in each block

$$\hat{q}_p^{t+1} = \sum_{q=1}^L \frac{\alpha_q J_{qp}}{\Delta + \sum_{r=0}^L J_{qr} E_r^t}, \quad (3) \\ E_p^{t+1} = \rho - \frac{\rho^2 \hat{q}_p^{t+1}}{\hat{q}_p^{t+1} + 1} G(\rho, \hat{q}_p^{t+1}),$$

where

$$G(\rho, \hat{q}) = \int \frac{dz}{\sqrt{2\pi}} \frac{z^2 e^{-\frac{z^2}{2}}}{\rho + (1 - \rho) e^{-\frac{z^2 \hat{q}}{2}} \sqrt{\hat{q} + 1}}. \quad (4)$$

We hence study the behavior of these state evolution equations. The only initialization that is algorithmically possible in practice is $E_p = \rho$ (or larger) for every block p . This is the main reason why we need the seed to exist - its higher undersampling ratio α_s has to ensure that at least the first w_s blocks will get reconstructed and the interaction of range w then has to ensure this reconstruction will propagate into the whole system.

III. THE FINITE- w PHASE TRANSITION

In the limit of many blocks $L \rightarrow \infty$, given that the seed remains of finite size, one has that the effective measurement ratio is equal to the bulk one, $\lim_{L \rightarrow \infty} \alpha_{\text{eff}} = \alpha_b$. This, combined with the limit $w \rightarrow \infty$ taken after $L \rightarrow \infty$ is used in the proofs of threshold saturation, i.e. optimal recovery of the signal is achieved down to α_c . In practical implementations, the interaction range w has to be relatively small since $w < L < N$, and also as we will see since the seed cost grows with w . In this Section we show, by analyzing numerically the state evolution equations, that if w is finite reconstruction is possible only down to some $\alpha_w > \alpha_c$ even when $L \rightarrow \infty$. The seed we use for this analysis is very large and very strong $w_s \gg 1$, $\alpha_s > 1$, the number of blocks L is large, and a flat interaction $g(x) = 1/2$ is used. The results then depend only on the parameters α_b , ρ , Δ , and w .

In Fig. 2 we draw the α_w transition line (above which reconstruction with spatially coupled matrix is achieved) for measurement noise $\Delta = 10^{-12}$ and different values of the interaction range from $w = 1$ to $w = 4$. We notice that for




Fig. 2. A phase diagram for the coupled system with $\Delta = 10^{-12}$, $L = 240$. The black solid line (uppermost) and the blue solid line (lowermost) mark, respectively, the BP threshold and the reconstructability limits of the single system. The dashed lines represent the reconstruction thresholds of the coupled system for different values of the coupling range. Increasing the coupling range, the threshold of the coupled system approaches the critical line of the single system maintaining a smaller and smaller bump at small ρ . Note that the Donoho-Tanner phase transition [17] well known in compressed sensing literature would be considerably above the α_{BP} line.

low values of ρ , the transition line of the coupled system α_w is well separated from the critical line (reconstructability limit) α_c of the single system. Whereas for larger values of ρ the threshold α_w is extremely close to α_c . The values of α_c used here are computed with the equations resulting from the replica method calculation from [10].

We are not aware of previous work where this observation of two separate regimes (large versus small gap from optimality at small versus large values of sparsity ρ) existing for small w would have been made. Interestingly, careful examination of the data presented in Fig. 2 of [9] leads to a conclusion that also in that case (also $w = 1$) this effect is visible. With growing w , instead, the transition line tends to lean over the critical line of the single system, until they will be identical for $w \gg 1$.

In Fig. 3 we give the same diagram as in Fig. 2, but this time at fixed $w = 1$ and for different values of the noise variance Δ . In this case the line α_w seems to have low- ρ part that is roughly Δ -independent (similarly to the α_{BP} threshold), and a high- ρ part that is always very close to the (Δ -dependent) reconstructability limit α_c .

For the coupled system, we can write the potential (negative free energy in physics) as a function of the error profile $\{E_p\}$ [10], the fixed points of the state evolution being given by the stationary points of the potential. As explained in [18] for a simpler system, namely the Curie-Weiss model, we can also obtain the free energy $\Phi(E)$ as a function of the mean-squared error by extremizing the “profile” free energy subject to the constraint $E = \sum_r E_r/L$. In [18] the authors also show that the potential function $\Phi(E)$ of the spatially coupled system in the ordered limit $L \rightarrow \infty$, $w \rightarrow \infty$ becomes the convex envelope of the free energy of the single block. Exactly




Fig. 3. A phase diagram for the coupled system with $w = 1$, $L = 240$. The black solid line (uppermost) marks the BP threshold of the single system (interestingly this line does not change visibly for the considered values of Δ). The solid and dashed lines in other colors represent, respectively, the critical line of the single system and the threshold for the coupled system for different noise levels.

at the critical ratio $\alpha_b = \alpha_c$ the two maxima of the potential function corresponding to optimal reconstruction and failed reconstruction respectively, are connected by a *plateau*. This means that, for any α_b infinitesimally larger than α_c , there is only one maximum of the potential, and it gives the MSE of the optimal reconstruction. This is the explanation (in terms of the potential) of the fact that the threshold line for the coupled system tends to the critical line of the single system when $w \rightarrow \infty$.

The way in which the potential behaves at finite w (i.e. the way in which it converges to the convex envelope) determines the structure of the threshold lines in Figs. 2 and 3. When the system is coupled with growing w , the potential of the single system tends to be flattened in the region in between the two fixed points and it is “dressed” with oscillations whose frequency increases with L and whose amplitude decreases with both L and w . In our case we see that the convergence is faster for larger ρ and is much slower for small ρ . When the finite- w threshold line is very close to the critical line of the single system the scenario is very similar to the one presented in [18], [19] for the Curie-Weiss model. On the other hand, when the separation between α_w and α_c is large, as it is for $w = 1$, low ρ and small Δ , convergence of $\alpha_w \rightarrow \alpha_c$ becomes much slower in w . The fundamental reason behind this remains to be unveiled in future work.

IV. THE SEEDING CONDITION

What we presented in the previous Section is valid for very large seed w_s and very large seeding ratio α_s . The seed has to be present in order to trigger the propagation of the reconstruction wave. However, given a fixed bulk-measurement-ratio $\alpha_b \in [\alpha_w, \alpha_{BP}]$, the effective measurement ratio (average over the whole system) α_{eff} will be slightly bigger than α_b and therefore for finite L we will not be able




Fig. 4. [Main frame] The wave propagation/non-propagation diagram as a function of the seed size w_s and seed strength α_s for $\rho = 0.4$, $\Delta = 10^{-12}$, $\alpha_b = 0.5$ and $L = 400$ for three different values of the interaction range $w = 1, 2, 10$. Below the lines the reconstruction wave does not propagate while above it does. [Inset] The effective measurement ratio along the transition line for the same values of the interaction range.

to access values of α_{eff} exactly down to α_w .

As an example, in Fig. 4 we show the propagation/non-propagation phase diagram as a function of the seed size and strength for $w = 1, 2, 10$, $\rho = 0.4$, $\Delta = 10^{-12}$, $\alpha_b = 0.5$ and $L = 400$. Just as it happens in the case of the spatially-coupled Curie-Weiss model [19], the transition line approaches a smooth limiting curve when the interaction range w grows. In the inset we show the effective measurement ratio α_{eff} along the transition line, i.e. the best we can do given the bulk value α_b , the interaction range w and the fixed number of blocks. Again, as in the Curie-Weiss model [19], the best choice (the one that minimizes α_{eff}) is $w = 1$. Notice that for $\rho = 0.4$, $\Delta = 10^{-12}$ and $w = 1$, we have $\alpha_w = 0.4619$, so the value of α_b we have chosen is above the threshold line for unitary interaction range. In other cases, as for example $\rho = 0.2$, $\Delta = 10^{-12}$ and $\alpha_b = 0.26$, the wave does not propagate with $w = 1$ for any choice of the seed as we can see from the diagram of Fig. 2 and the best choice becomes therefore $w = 2$.

V. THE ROLE OF THE INTERACTION SHAPE

Until now we have considered the case of a constant interaction $g(x) = 1/2$ on the interval $[-1, 1]$. On the other hand, in our model any properly normalized function is allowed. A first simple generalization is an interaction of the form

$$g(x) = \frac{1}{2} + Ax \quad (5)$$

with $A \in [-1/2, 1/2]$.

The value of the constant A influences three different quantities, namely:

- the position of the finite- w transition α_w
- the speed v of the propagating wave defined as the number of reconstructed blocks per iteration




Fig. 5. The propagation speed in the interval $\alpha_b \in [\alpha_c, \alpha_{BP}]$ for $\rho = 0.12$ and $\Delta = 10^{-12}$, $w = 3$, $L = 400$. We report the speed for three different angular coefficients of the interaction function $g(x)$, namely $A = -0.5, 0, 0.5$. The two vertical lines mark α_c (left) and α_{BP} (right). Note that the transition α_w (when the speed gets positive) depends considerably on A .

- the necessary size and strength of the seed

All the three characteristics are in favor of a negative A coefficient, meaning that each block interacts more strongly backwards than forwards. In this case, in fact, the finite- w transition moves closer to α_c , the propagation speed is increased and the necessary seeding is weaker. Viceversa, when the angular coefficient is positive, the propagation speed is lower and the seeding conditions must be stronger.

Note that in the original work on spatial coupling in compressed sensing [9], [10] the interaction was chosen strongly asymmetric, our present study confirms this choice as the more efficient one.

In Fig. 5 we show the propagation speed in the interval $\alpha_b \in [\alpha_c, \alpha_{BP}]$ for three different values of A , while in Fig. 6 we report the seeding diagrams for several values of the angular coefficient.

VI. CONCLUSION

In this work we have studied the design of spatially coupled measurement matrices for compressed sensing. In particular we have chosen the structure defined in Fig. 1, and evaluated its efficiency as a function of the interaction range w , the interaction shape $g(x)$, and the size and strength of the seed w_s and α_s . We find that, for sufficiently large seed, there is a threshold line α_w for the coupled system below which reconstruction fails. This line lies in between the lowest possible subsampling rate at which Bayes-optimal inference succeeds α_c and the α_{BP} threshold of the single system. The line α_w is considerably separated from α_c when the interaction range w and density ρ are both sufficiently small. We also studied the optimality of the seeding condition, concluding that for small termination cost small range w is better. Finally we analyzed the speed of the reconstruction as a function of the interaction shape, given by the function $g(x)$ and concluded that a strong asymmetry is advantageous.




Fig. 6. The propagation phase diagram as a function of the seed features for different values of the angular coefficient in the interaction function $g(x)$, with $\rho = 0.4$, $\Delta = 10^{-12}$, $\alpha_b = 0.5$, $w = 1$ and $L = 400$.

ACKNOWLEDGMENT

This work has been supported in part by the ERC under the European Unions 7th Framework Programme Grant Agreement 307087-SPARCS, and by the project TASC of the Labex PALM.

REFERENCES

- [1] A. Jimenez Felstrom and K. Zigangirov, "Time-varying periodic convolutional codes with low-density parity-check matrix," *IEEE Transactions on Information Theory*, vol. 45, no. 6, pp. 2181–2191, 1999.
- [2] M. Lentmaier, A. Sridharan, K. S. Zigangirov, and D. Costello, "Terminated ldpc convolutional codes with thresholds close to capacity," in *Information Theory, 2005. ISIT 2005. Proceedings. International Symposium on*. IEEE, 2005, pp. 1372–1376.
- [3] M. Lentmaier, A. Sridharan, D. J. Costello, and K. S. Zigangirov, "Iterative decoding threshold analysis for ldpc convolutional codes," *IEEE Transactions on Information Theory*, vol. 56, no. 10, pp. 5274–5289, 2010.
- [4] S. Kudekar, T. J. Richardson, and R. L. Urbanke, "Threshold saturation via spatial coupling: Why convolutional ldpc ensembles perform so well over the bec," *IEEE Transactions on Information Theory*, vol. 57, no. 2, pp. 803–834, 2011.
- [5] S. Kudekar, T. Richardson, and R. Urbanke, "Spatially coupled ensembles universally achieve capacity under belief propagation," in *Information Theory Proceedings (ISIT), 2012 IEEE International Symposium on*. IEEE, 2012, pp. 453–457.
- [6] D. L. Donoho, "Compressed sensing," *IEEE Trans. Inform. Theory*, vol. 52, p. 1289, 2006.
- [7] E. J. Candès and M. B. Wakin, "An introduction to compressive sampling," *IEEE Signal Processing Magazine*, vol. 25, no. 2, pp. 21–30, Mar. 2008.
- [8] S. Kudekar and H. Pfister, "The effect of spatial coupling on compressive sensing," in *Communication, Control, and Computing (Allerton)*, 2010, pp. 347–353.
- [9] F. Krzakala, M. Mézard, F. Sausset, Y. Sun, and L. Zdeborová, "Statistical physics-based reconstruction in compressed sensing," *Phys. Rev. X*, vol. 2, p. 021005, 2012.
- [10] F. Krzakala, M. Mézard, F. Sausset, Y. Sun, and L. Zdeborová, "Probabilistic reconstruction in compressed sensing: Algorithms, phase diagrams, and threshold achieving matrices," *J. Stat. Mech.*, p. P08009, 2012.
- [11] D. L. Donoho, A. Javanmard, and A. Montanari, "Information-theoretically optimal compressed sensing via spatial coupling and approximate message passing," in *Proc. of the IEEE Int. Symposium on Information Theory (ISIT)*, 2012.
- [12] A. Yedla, Y.-Y. Jian, P. S. Nguyen, and H. D. Pfister, "A simple proof of threshold saturation for coupled scalar recursions," in *7th International Symposium on Turbo Codes and Iterative Information Processing (ISTC)*. IEEE, 2012, pp. 51–55.
- [13] D. L. Donoho, A. Maleki, and A. Montanari, "Message-passing algorithms for compressed sensing," *Proc. Natl. Acad. Sci.*, vol. 106, no. 45, pp. 18 914–18 919, 2009.
- [14] S. Rangan, "Generalized approximate message passing for estimation with random linear mixing," in *IEEE International Symposium on Information Theory Proceedings (ISIT)*, 2011, pp. 2168–2172.
- [15] M. Bayati, M. Lelarge, and A. Montanari, "Universality in polytope phase transitions and message passing algorithms," *arXiv preprint arXiv:1207.7321*, 2012.
- [16] S. Kudekar, T. Richardson, and R. Urbanke, "Wave-like solutions of general one-dimensional spatially coupled systems," *arXiv preprint arXiv:1208.5273*, 2012.
- [17] D. L. Donoho and J. Tanner, "Sparse nonnegative solution of underdetermined linear equations by linear programming," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 102, no. 27, pp. 9446–9451, 2005.
- [18] S. H. Hassani, N. Macris, and R. Urbanke, "Chains of mean-field models," *Journal of Statistical Mechanics: Theory and Experiment*, vol. 2012, no. 02, p. P02011, 2012.
- [19] F. Caltagirone, S. Franz, R. Morris, and L. Zdeborová, "Dynamics and termination cost of spatially coupled mean-field models," *arXiv:1310.2121v1*, 2013.